

METHOD AND APPARATUS FOR SOLVING AN OPTIMIZATION PROBLEM

CLAIM OF BENEFIT TO PRIOR APPLICATION

5 This patent application claims the benefit of U.S. Provisional Patent Application 60/427,131, filed 11/18/02.

FIELD OF THE INVENTION

The invention is directed towards a method and apparatus for solving an optimization problem.

BACKGROUND OF THE INVENTION

10

Design engineers design IC's by transforming logical or circuit descriptions of IC's into geometric descriptions, called layouts. IC layouts typically include circuit modules with pins, and interconnect lines that connect the circuit-module pins. A net is typically defined as a collection of pins that need to be connected. A list of all or some of the nets in a layout is referred to as a net list.

15

To create layouts, design engineers typically use electronic design automation ("EDA") applications. These applications provide sets of computer-based tools for creating, editing, and analyzing IC design layouts. Routers are one type of EDA tool. A router defines routes for interconnect lines that connect the pins of nets.

20

A router typically receives a routing problem that requires it to identify a route for each net in a net list. Some routers solve such a routing problem by first identifying multiple routes for each net. These routers then try to select a combination of routes that includes one route for each net and that optimizes a particular objective function. The objective function

BEST AVAILABLE COPY

THIS PAGE IS BLANK

typically expresses one or more metric costs, such as length of routes, congestion in the design, etc.

Routers of this sort use different techniques to select the combination of routes from the set of identified solutions. One technique is randomized rounding. Generally, this
5 technique is a probabilistic method that converts an exact solution of a relaxed problem into an approximate solution to an original problem. With randomized rounding, an EDA router interprets fractional routing solutions (*e.g.*, solutions provided by a linear program) as probabilities for rounding the solutions. Numerous examples of this technique can be found in the literature. One such reference is disclosed in *Randomized Algorithms*, by Rajeev
10 Motwani and Prabhakar Raghavan, Cambridge University Press (1995, 1997).

Randomized rounding works well in certain situations but not in others. For instance, one type of routing is global routing. Global routing typically identifies routes between cells, which partition the layout into several regions. Such cells are called Gcells. In global routing, randomized rounding works well when the Gcells are large, but does not work as well when
15 the Gcells are small. Therefore, there is a need in the art for a better method for selecting a combination of routes from a set of routes that includes one or more routes for each net in a net list. More generally, there is a need for a better method for solving optimization problems.

BEST AVAILABLE COPY

THIS PAGE IS BLANK

SUMMARY OF THE INVENTION

Some embodiments of the invention provide a method of solving an optimization problem. The problem includes a plurality of elements, and one or more solutions have been previously identified for each element. The method specifies a first solution set that has one
5 identified solution for each element. In some embodiments, the method then iteratively examines all the elements of the problem. During the examination of each particular element, the method iteratively examines all the identified solutions for the particular element. During the examination of each particular solution, the method replaces the current solution for the particular element in the first solution set with a previously unexamined solution for the
10 particular element if the replacement would improve the first set.

BEST AVAILABLE COPY

THIS PAGE IS BLANK

BRIEF DESCRIPTION OF THE DRAWINGS

The novel features of the invention are set forth in the appended claims. However, for the purpose of explanation, several embodiments of the invention are set forth in the following figures.

5 **Figure 1** illustrates a process of some embodiments of the invention.

Figures 2-4 illustrate several simple examples of Gcells and global routes.

Figure 5 illustrates a process for specifying a group of routes for a set of nets.

Figure 6 conceptually illustrates a computer system with which one embodiment of the invention is implemented.

10

BEST AVAILABLE COPY

THIS PAGE IS BLANK

DETAILED DESCRIPTION OF THE INVENTION

In the following description, numerous details are set forth for the purpose of explanation. However, one of ordinary skill in the art will realize that the invention may be practiced without the use of these specific details. In other instances, well-known structures and devices are shown in block diagram form in order not to obscure the description of the invention with unnecessary detail.

Figure 1 illustrates a process 100 of some embodiments of the invention. This process solves a global routing optimization problem. This problem requires the process to identify a global route for each net in a netlist. The netlist includes some or all the nets in a region of a design layout. One of ordinary skill will realize that other embodiments might solve other types of optimization problems, such as integer flow and transportation, graph covering and coloring, maximum logic clause satisfaction, etc.

The process 100 solves the routing problem by first identifying several global routes for each net and then exploring the solution space to find an optimal solution. The process 100 can quickly and flexibly identify a good combination of routes for a set of nets. To explore the solution space, the process has three nested loops: an outer loop, a middle loop, and an inner loop. Generally (1) each iteration of the outer loop (which starts at 110) explores the solution space from a different starting group of routes, (2) for a particular starting group of routes, each iteration of the middle loop (which starts at 115) explores the solution space for a particular order of the nets, and (3) for each particular order of the nets, each iteration of the inner loop (which starts at 120) iteratively examines the routes of different nets. These loops are further described below.

The process 100 initially identifies (at 105) several sets of global routes. Each set has one global route for each net in the net list. Different embodiments use different techniques to

BEST AVAILABLE COPY

THIS PAGE IS BLANK

identify these sets. For instance, some embodiments use techniques like those described in United States Patent Application 10/013,819, filed on December 7, 2001. Other embodiments use techniques like those described in United States Patent Application entitled "Method and Apparatus for Routing", filed concurrently with the present application with the Express Mail
5 Number EV169571570US, and a serial number **. This concurrently filed application is incorporated herein by reference.

Each identified global route connects several Gcells. Some or all of the Gcells connected by a global route for a net contain routable elements (*e.g.*, pins) of the net. **Figures 2-4** illustrate simple examples of Gcells and global routes. Specifically, **Figure 2** illustrates a
10 4x4 partitioning grid that partitions a design-layout region into 16 Gcells, which are numbered as Gcells 0-15. Each of these Gcells can be connected to its neighboring Gcells through one of twenty-four horizontal or vertical edges E0-E23 (illustrated in **Figure 2**) between the Gcells. **Figures 3** and **4** illustrate two global routes for a net that has routable elements in Gcells 1 and 8. Both the routes are equal-length spanning trees. United States
15 Application 10/013,819 discloses other examples of global routes, including ones with diagonal edges between Gcells. The above-incorporated application also discloses other examples of global routes, Gcells, and diagonal edges.

After identifying several sets of global routes at 105, the process specifies (at 110) a current group of routes. The specified group includes a route for each net in the net list. In
20 some embodiments, all the routes in the specified group are from one set of global routes that the process identified at 105. In other embodiments, the routes of the specified group are from two or more of the sets identified at 105. For instance, if the netlist included eight nets, the specified group of routes might include five routes from one identified set and three

BEST AVAILABLE COPY

THIS PAGE IS BLANK

routes from another identified set. **Figure 5** illustrates how one such group of routes can be specified. This figure is further described below.

After 110, the process specifies (at 115) an order for the nets. Different embodiments use different techniques to specify such an order. For instance, some embodiments specify an order based on descending entropies of the nets, others based on Steiner tree lengths of the
5 nets, and yet others based on bounding boxes of the nets. Still other embodiments randomly order the nets at 115.

The process 100 iterates through 115 several times for a particular starting group of routes that it selects at 110. In one or more of the iterations through 115, the process might
10 specify an order for the nets that is based on the frequency of occurrence of the routes for the nets in the identified set of routes. Specifically, in some embodiments, a net can have the same route in two or more of the sets identified at 105. Accordingly, a frequency parameter (*e.g.*, a probability) can be specified for each identified route of each net. For instance, the process might specify eight sets of routes at 105. However, for a particular net, the eight sets
15 might only include three different routes, a first that is in five sets, a second that is in two sets, and a third that is in only one set. Based on the times that the first, second, and third routes appear in the identified sets, probability values $5/8$, $1/4$, and $1/8$ can be respectively assigned to the first, second, and third routes.

One order that can be specified in an iteration through 115 is an order that is based on
20 descending maximum probability values of the nets. For instance, in a simple example, the net list might include three nets, with the first net having a route that appears in 90% of the identified solutions (specified at 105), the second net having a route that appears in 70% of the identified solutions, and the third net having a route that appears in 55% of the identified solutions. In this example, an order that is based on descending maximum probability values

BEST AVAILABLE COPY

THIS PAGE IS BLANK

would specify the first net, followed by the second net, followed by the third net. One of ordinary skill will realize that other orders can be specified that are derived from probability values relating to frequency of routes in the identified sets of routes.

After 115, the process 100 selects (at 120) a net according to the order specified at 115. For the net selected at 120, the process then determines (at 122) whether the identified sets of routes include at least two unique routes for the selected net. If not, the process transitions to 145, which is further described below.

Otherwise, the process selects (at 125) a route that is one of the identified routes for the selected net but is not the current route for this net in the current group of routes. The process next determines (at 130) whether replacing the selected net's route in the current group with the route selected at 125 would improve the quality of the routing. This determination entails computing two metric scores. One metric score quantifies the quality of the current group of routes. The other metric score quantifies the quality of the current group of routes with the selected net's route in the current group replaced with the route selected at 125.

Different embodiments of the invention use different objective functions to compute a metric score. Some embodiments use the following function (A), which has two components, a first (congestion) component and a second (length) component.

$$\text{Function} = \sum_{j=1}^m Y_j e^{\epsilon_1 \cdot \frac{\text{usage}(j)}{\text{goal}(\text{layer}(j)) \cdot \text{capacity}(j)}} + Y_1 \sum_{n=1}^p e^{\epsilon_2 \cdot \frac{\text{length}(n)}{\text{lower_bound}(n)}} \quad (\text{A})$$

In this function, (1) j represents one of m edges between Gcells, (2) e is the base of the natural logarithm, (3) ϵ_1 and ϵ_2 are user-adjustable parameters, (4) $\text{usage}(j)$ is the number of routes in the group of routes that use the edge j , (5) $\text{capacity}(j)$ is the estimated maximum number of available tracks along the edge j , (6) $\text{goal}(\text{layer}(j))$ is a target upper bound on the

BEST AVAILABLE COPY

THIS PAGE IS BLANK

congestion ratio on the design-layout layer that contains edge j , (7) n is one of the p nets in the netlist, (8) $\text{length}(n)$ is the length of the route of net n in the group of nets, (9) $\text{lower_bound}(n)$ is a lower bound for the route length of the net n , and (10) Y_j and Y_l are constants that normalize the length and congestion components. When this function is used, a smaller metric score means better quality.

In some embodiments, an edge j in function (A) can be a planar or non-planar edge. A non-planar edge represents a via, which typically is a transition between two adjacent layers within a Gcell. Accordingly, in these embodiments, the congestion component of function (A) not only accounts for planar congestion, but also accounts for via (non-planar) congestion. Instead of merging the via congestion cost into the overall congestion cost, the function (A) in other embodiments can have a third component that expresses the via congestion cost. Having a third component would allow the user-adjustable epsilon parameter ϵ for via congestion to be different than the user-adjustable epsilon parameter for the planar wire congestion.

If the quality (*i.e.*, the metric score) of the current group of routes does not improve with the route selected at 125, the process transitions from 130 to 140, which is further described below. On the other hand, if the quality of the current group of routes improves with the route selected at 125, the process replaces (at 135) the selected net's current route in the group with the route selected at 125, and then transitions to 140.

At 140, the process determines whether, during its current loop through 120-145 (*i.e.*, during the current instance of the inner loop), it has examined all the routes that it identified at 105 for the selected net. If not, the process transitions back to 125 to select, for the selected net, a route that it has not yet examined in its current loop through 120-145.

BEST AVAILABLE COPY

THIS PAGE IS BLANK

Otherwise, the process determines (at 145) whether it has examined all the nets in the netlist during its current loop through 115-150 (*i.e.*, during the current instance of the middle loop). If not, the process transitions back to 120 to select another net according to the order specified at 115.

5 When the process determines (at 145) that it has examined all the nets in the netlist during its current loop through 115-150, it determines (at 150) whether it has examined enough permutations of routes based on the same initial group of routes that were specified at 110. Different embodiments base the determination at 150 on different criteria. Examples of such criteria include the number of iterations, the percentage of improvement in the routing
10 quality between the last two or more iterations, etc. Some embodiments might base this determination on different criteria during different iterations through 150.

 If the process determines (at 150) that it has not examined enough permutations, the process transitions back to 115 to specify an order for the nets. In some embodiments, this new order is different from all the orders that the process set in its previous iterations through
15 115, while, in other embodiments, this order might be the same as one of the previously set ones. For the newly specified order, the process then performs operations 120-150 as described above.

 Once the process 100 determines (at 150) that it has done enough iterations through the nets for the starting solution specified at 110, the process compares (at 155) the metric
20 score of the current group of routes with a Best_Score that represents the score of the best group of routes that the process has identified up to this stage. If the current group's score is better than the Best_Score, the process stores (at 155) the current group of routes as the Best_Solution that it has encountered up to this stage. At the start of the process 100, some embodiments initialize the Best_Score to a very large number.

BEST AVAILABLE COPY

THIS PAGE IS BLANK

Next, the process determines (at 160) whether it has started its exploration of the solution space from a sufficient number of different starting points. If not, the process transitions back to 110 to specify another initial group of routes as the current group of routes. In some embodiments, this group of routes is different from all the groups that the process previously specified in its previous iteration through 110. For the newly specified group, the process then performs operations 115-160 as described above. When the process determines (at 160) that it has examined enough starting points for its search, the process identifies (at 165) the Best_Solution as the solution to the routing problem. The process then terminates.

Figure 5 illustrates a process 500 that specifies a group of routes for a set of nets after a set of routes has been defined for each net in the set. The set of nets includes at least two nets. Also, in some embodiments, the set of routes for each net includes at least one route. Some embodiments use this process to define one of the starting groups of routes at 110 of the process 100.

The process 500 iteratively selects one identified route for each net. Specifically, during each iteration, the process selects the identified route that least increases an exponential tracking cost, as further described below. The process 500 initially orders the nets based on decreasing maximum probability values of the routes for the nets. Such probability values and orders based on such values were described above in connection with operation 115 of process 100.

After 505, the process 500 selects (at 510) the first net according to the order specified at 505. It then selects (at 515) the route for the first net that has the best metric score. In some embodiments, the process uses the following objective function (B) to compute the metric scores of each route for the first net.

BEST AVAILABLE COPY

THIS PAGE IS BLANK

$$\text{Function} = Y_1 * e^{\epsilon_2 * \frac{\text{length}(\text{net})}{\text{lower_bound}(\text{net})}} + \sum_{j=1}^m Y_j e^{\epsilon_1 * \frac{\text{usage}(j)}{\text{goal}(\text{layer}(j)) * \text{capacity}(j)}} \quad (\text{B})$$

In this function, (1) j represents one of m edges between Gcells, (2) e is the base of the natural logarithm, (3) ϵ_1 and ϵ_2 are user-adjustable parameters, (4) $\text{usage}(j)$ indicates whether the route uses the edge j , (5) $\text{capacity}(j)$ is the estimated maximum number of available tracks along the edge j , (6) $\text{goal}(\text{layer}(j))$ is a target upper bound on the congestion ratio on the design-layout layer that contains edge j , (7) net is the selected first net, (8) $\text{length}(\text{net})$ is the length of the route for the net, (9) $\text{lower_bound}(\text{net})$ is the lower bound for the route length of the net, and (10) Y_j and Y_1 are constants that normalize the length and congestion components.

The process then selects (at 520) the next net according to the order specified at 505. From the set of routes identified for the selected net, the process then selects (at 525) the route that results in the smallest exponential tracking metric score. Some embodiments use the following function (C), which is similar to the above-described function (A).

$$\text{Function} = \sum_{j=1}^m Y_j e^{\epsilon_1 * \frac{\text{usage}(j)}{\text{goal}(\text{layer}(j)) * \text{capacity}(j)}} + Y_1 \sum_{n=1}^p e^{\epsilon_2 * \frac{\text{length}(n)}{\text{lower_bound}(n)}} \quad (\text{C})$$

In this function, (1) j represents one of m edges between Gcells, (2) e is the base of the natural logarithm, (3) ϵ_1 and ϵ_2 are user-adjustable parameters, (4) $\text{usage}(j)$ is the number of routes that have been selected thus far that use the edge j , (5) $\text{capacity}(j)$ is the estimated maximum number of available tracks along the edge j , (6) $\text{goal}(\text{layer}(j))$ is a target upper bound on the congestion ratio on the design-layout layer that contains edge j , (7) n is one of the p nets selected thus far, (8) $\text{length}(n)$ is the length of the route of net n , (9) $\text{lower_bound}(n)$ is a lower bound for the route length of the net n , and (10) Y_j and Y_1 are constants that normalize the length and congestion components.

BEST AVAILABLE COPY

THIS PAGE IS BLANK

To select (at 525) the route for the selected net, some embodiments compute the value of function (C) for each route in the set of routes identified for the selected net. For each particular route of the selected net, the function (C) is computed for a set of routes that includes the particular route and all routes previously selected (at 510 or 525) for previously
5 selected nets.

After 525, the process determines (at 530) whether it has selected a route for each net. If not, the process selects (at 520) the next net according to the specified order, and then selects (at 525) the route (from the set of routes identified for the selected net) that results in the smallest exponential tracking metric score. When the process determines at 530 that it has
10 selected a route for each net, it terminates.

Figure 6 presents a computer system with which one embodiment of the present invention is implemented. Computer system 600 includes a bus 605, a processor 610, a system memory 615, a read-only memory 620, a permanent storage device 625, input devices 630, and output devices 635.

15 The bus 605 collectively represents all system, peripheral, and chipset buses that support communication among internal devices of the computer system 600. For instance, the bus 605 communicatively connects the processor 610 with the read-only memory 620, the system memory 615, and the permanent storage device 625.

From these various memory units, the processor 610 retrieves instructions to execute
20 and data to process in order to execute the processes of the invention. The read-only-memory (ROM) 620 stores static data and instructions that are needed by the processor 610 and other modules of the computer system. The permanent storage device 625, on the other hand, is a read-and-write memory device. This device is a non-volatile memory unit that stores instruction and data even when the computer system 600 is off. Some embodiments of the

BEST AVAILABLE COPY

THIS PAGE IS BLANK

invention use a mass-storage device (such as a magnetic or optical disk and its corresponding disk drive) as the permanent storage device 625. Other embodiments use a removable storage device (such as a floppy disk or zip® disk, and its corresponding disk drive) as the permanent storage device.

5 Like the permanent storage device 625, the system memory 615 is a read-and-write memory device. However, unlike storage device 625, the system memory is a volatile read-and-write memory, such as a random access memory. The system memory stores some of the instructions and data that the processor needs at runtime. In some embodiments, the invention's processes are stored in the system memory 615, the permanent storage device
10 625, and/or the read-only memory 620.

 The bus 605 also connects to the input and output devices 630 and 635. The input devices enable the user to communicate information and select commands to the computer system. The input devices 630 include alphanumeric keyboards and cursor-controllers. The output devices 635 display images generated by the computer system. For instance, these
15 devices display IC design layouts. The output devices include printers and display devices, such as cathode ray tubes (CRT) or liquid crystal displays (LCD).

 Finally, as shown in **Figure 6**, bus 605 also couples computer 600 to a network 665 through a network adapter (not shown). In this manner, the computer can be a part of a network of computers (such as a local area network ("LAN"), a wide area network ("WAN"),
20 or an Intranet) or a network of networks (such as the Internet). Any or all of the components of computer system 600 may be used in conjunction with the invention. However, one of ordinary skill in the art would appreciate that any other system configuration may also be used in conjunction with the present invention.

 While the invention has been described with reference to numerous specific details,

BEST AVAILABLE COPY

THIS PAGE IS BLANK

one of ordinary skill in the art will recognize that the invention can be embodied in other specific forms without departing from the spirit of the invention. For instance, some embodiments might use different cost functions than those described above. Several of the above-described cost functions divide the length of each net's route by a lower-bound for the route length of the net. Instead of the lower-bound, other embodiments might use as a divisor another estimate of route length for the net. For example, some embodiments might use the length of a congestion-unaware route for the net. The generation of a congestion-unaware route is described in the above-incorporated application.

Also, the cost function of some embodiments might include only congestion components and not include any wirelength components. Alternatively, some embodiments might use a different wirelength component, such as

$$\text{Wirelength Component} = Y_L * e^{\frac{\sum_{n=1}^p \text{Length}(n)}{\epsilon_L * \text{T.E.L.}}},$$

where T.E.L. is the total length of the estimated routes (e.g., congestion-unaware routes) for the p nets. The estimated lengths could be the lower-bound lengths or some other estimated lengths. Other embodiments might use an equation that uses both this wirelength component and the wirelength component of the above-described equations.

Still other embodiments might slightly modify the wirelength component of the above-described equations slightly. For instance, some embodiments described above express

the wirelength component as $Y_L \sum_{n=1}^p e^{\frac{\epsilon_L * \text{length}(n)}{\text{lower_bound}(n)}}$. Instead of this formulation, some

embodiments might use $Y_L \sum_{n=1}^p e^{\frac{\epsilon_L * \text{length}(n)}{B(n) * \text{lower_bound}(n)}}$, where B is a factor related to the importance of the net. This factor B is smaller (e.g., it is 1) for important critical nets (e.g., time critical nets) that need shorter routes, while it is larger (e.g., it is 3) for non-critical nets that can have

BEST AVAILABLE COPY

THIS PAGE IS BLANK

longer routes. Accordingly, this factor causes the selection of shorter routes for critical nets, by increasing the wirelength cost of these nets much faster than the wirelength cost of non-critical nets. One of ordinary skill will realize that the net-importance factor B can also be multiplied by the net's route length (*i.e.*, by $\text{length}(n)$). In this situation, the factor B is larger
5 for an important time-critical net, and is smaller for a non-critical net.

Several embodiments were described above for solving a global routing optimization problem. One of ordinary skill will realize, however, that other embodiments can be used to solve other classes of optimization problems. In addition, the process 100 initially identifies (at 105) sets of solutions, where each set includes a global route for each net in the net list.
10 Other embodiments, however, might not identify such sets. For instance, some embodiments might just identify (at 105) one or more routes for each net, and then specify (at 110) different starting groups of routes by selecting different routes for some or all of the nets. Thus, one of ordinary skill in the art would understand that the invention is not to be limited by the foregoing illustrative details, but rather is to be defined by the following claims.

15

BEST AVAILABLE COPY

THIS PAGE IS BLANK

CLAIMS

We claim:

1. A method of solving an optimization problem that includes a plurality of elements, wherein one or more solutions are identified for each element, the method
5 comprising:

a) specifying a first solution set that has one identified solution for each element;

b) selecting a first element;

c) in the first solution set, replacing the current solution for the first
10 element with another identified solution for the first element if the replacement would improve the first set.

2. The method of claim 1, wherein the problem is a routing problem, each element is a net in a design layout, and each solution is a route.

3. The method of claim 1, wherein the replacement would improve the first set if
15 the replacement would improve a metric score that is computed for the first set.

4. The method of claim 3 further comprising:

a) computing a first metric score for the first set;

b) computing a second metric score for the first set when the current solution in the first set is replaced by the other identified solution for the first element;

c) replacing the current solution with the other identified solution if the
20 second metric score is better than the first metric score.

5. The method of claim 4, wherein the metric scores are computed based on an objective function.

BEST AVAILABLE COPY

THIS PAGE IS BLANK

6. The method of claim 5, wherein a lower metric score is better than a higher metric score.

7. The method of claim 1 further comprising:
identifying a plurality of solution sets for the problem, wherein each
5 solution set includes a solution for each element;

8. The method of claim 1, wherein the first set is one of the identified solution sets.

9. The method of claim 1, wherein the first set includes solutions for different elements that are taken from different identified solution sets.

10 10. The method of claim 1 further comprising:
iteratively examining all the identified solutions for the first element, during
each examination of a particular identified solution for the first element, replacing the current
solution for the first element in the first set with the particular identified solution if the
replacement would improve the first set.

15 11. The method of claim 10 further comprising:
a) selecting a second element, the second element having an initial
solution in the first set;

b) iteratively examining all the identified solutions for the second
element, during each examination of a particular identified solution for the second element,
20 replacing the current solution for the second element in the first set with the particular
identified solution if the replacement would improve the first set.

12. The method of claim 1 further comprising:

a) iteratively examining all elements of the problem,

BEST AVAILABLE COPY

THIS PAGE IS BLANK

b) during the examination of each particular element, iteratively examining all the identified solutions for the particular element,

(i) during each examination of an identified solution for the particular element, replacing the current solution for the particular element in the first set with the particular identified solution if the replacement would improve the first set.

13. The method of claim 12, wherein each solution for a particular element has a computed occurrence frequency, wherein the elements are examined based on an order that depends on the occurrence frequencies of the identified solutions of the elements.

14. The method of claim 13, wherein the elements that have solutions with the higher occurrence frequencies are examined first.

15. The method of claim 12 further comprising:

a) iteratively specifying solution sets and attempting to improve the specified solution sets, wherein each solution set has one identified solution for each element;

b) during each iteration for a particular solution set, iteratively examining all elements of the problem;

i) during the examination of each particular element, iteratively examining all the identified solutions for the particular element,

during each examination of a particular identified solution for the particular element, replacing the current solution for the particular element in the particular solution set with the particular identified solution if the replacement would improve the particular solution set.

16. A computer readable medium that stores a computer program for solving an optimization problem that includes a plurality of elements, wherein one or more solutions are identified for each element, the computer program comprising sets of instructions for:

BEST AVAILABLE COPY

THIS PAGE IS BLANK

a) specifying a first solution set that has one identified solution for each element;

b) selecting a first element;

c) in the first solution set, replacing the current solution for the first element with another identified solution for the first element if the replacement would improve the first set.

17. The computer readable medium of claim 16, wherein the computer program further comprises sets of instructions for:

iteratively examining all the identified solutions for the first element;

replacing, during each iteration, the current solution for the first element in the first set with a previously unexamined solution for the first element if the replacement would improve the first set.

18. The computer readable medium of claim 17, wherein the computer program further comprises sets of instructions for:

a) iteratively examining all elements of the problem,

b) during the examination of each particular element,

(i) iteratively examining all the identified solutions for the particular element, and

(ii) during each iteration, replacing the current solution for the particular element in the first set with a previously unexamined solution for the particular element if the replacement would improve the first set.

19. The computer readable medium of claim 18, wherein the computer program further comprises sets of instructions for:

BEST AVAILABLE COPY

THIS PAGE IS BLANK

a) iteratively specifying solution sets and attempting to improve the specified solution sets, wherein each solution set that has one identified solution for each element;

b) iteratively examining, during each iteration for a particular solution set,
5 all elements of the problem;

i) during the examination of each particular element,

(1) iteratively examining all the identified solutions for the particular element, and

(2) during each iteration, replacing the current solution for the
10 particular element in the particular solution set with a previously unexamined solution for the particular element if the replacement would improve the particular solution set.

BEST AVAILABLE COPY

THIS PAGE IS BLANK

ABSTRACT

Some embodiments of the invention provide a method of solving an optimization problem. The problem includes a plurality of elements, and one or more solutions have been previously identified for each element. The method specifies a first solution set that has one identified solution for each element. In some embodiments, the method then iteratively examines all the elements of the problem. During the examination of each particular element, the method iteratively examines all the identified solutions for the particular element. During the examination of each particular solution, the method replaces the current solution for the particular element in the first solution set with a previously unexamined solution for the particular element if the replacement would improve the first set.

BEST AVAILABLE COPY

THIS PAGE IS BLANK

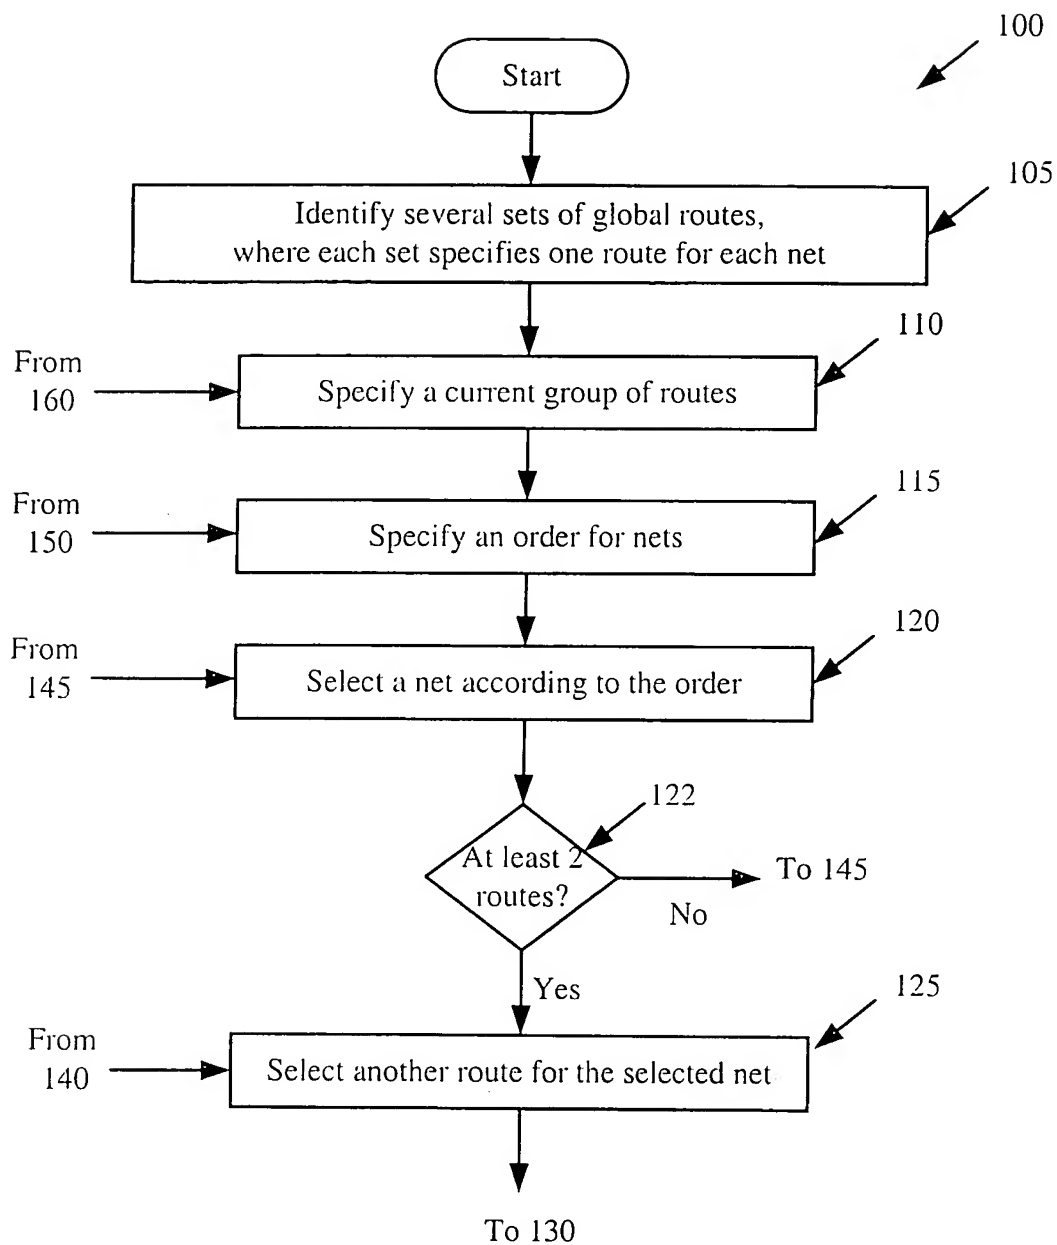


Figure 1A

Figure 1: Figure 1A
Figure 1B

BEST AVAILABLE COPY

THIS PAGE IS BLANK




Figure 1B

BEST AVAILABLE COPY

THIS PAGE IS BLANK




Figure 2




Figure 3




Figure 4

BEST AVAILABLE COPY

THIS PAGE IS BLANK

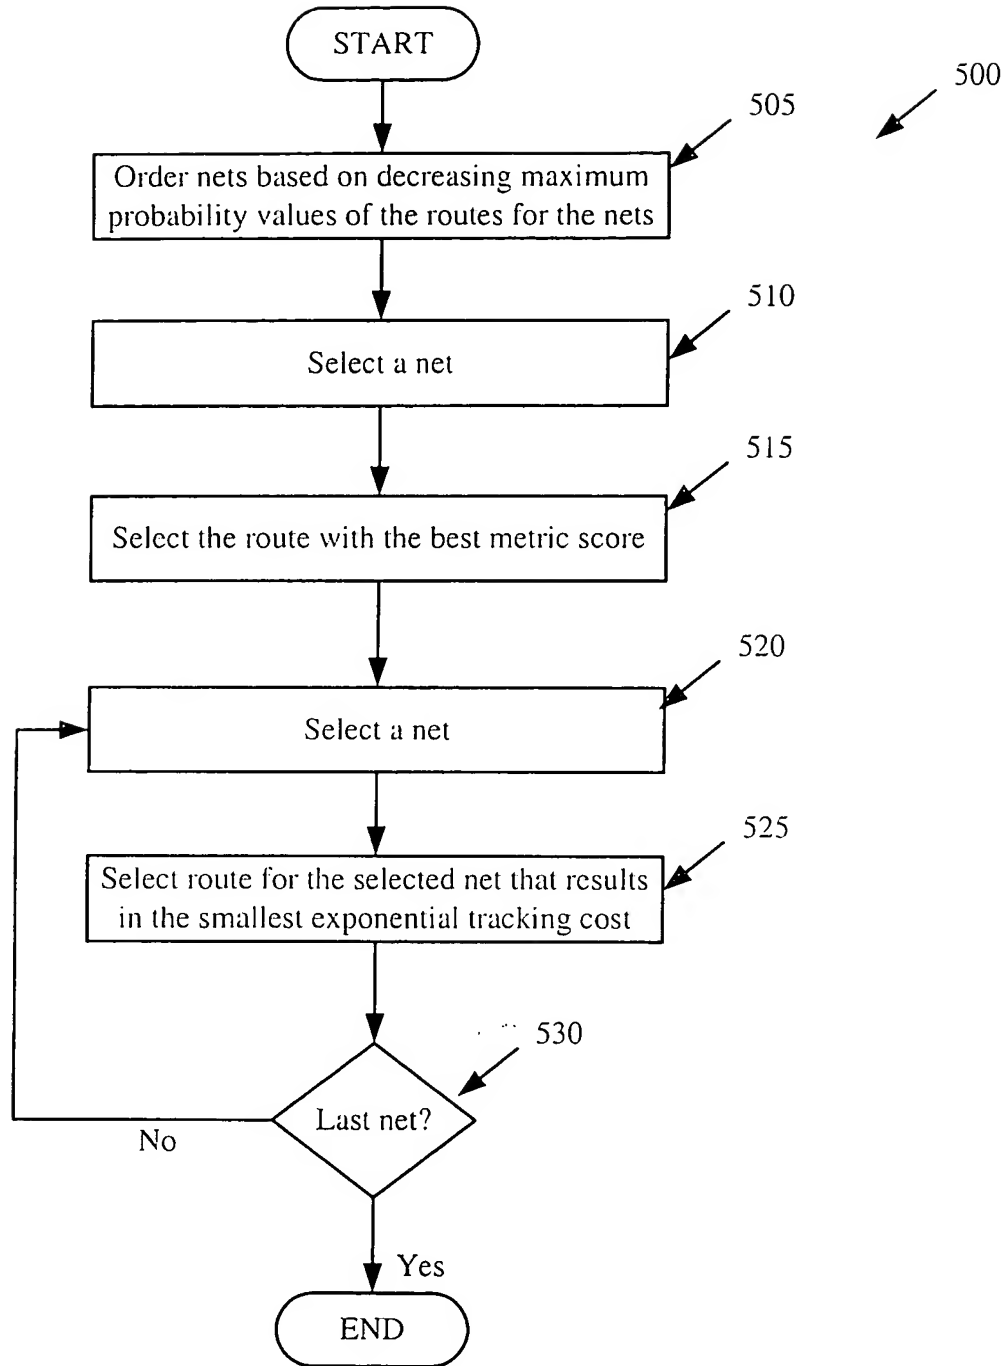


Figure 5

BEST AVAILABLE COPY

THIS PAGE IS BLANK




Figure 6

BEST AVAILABLE COPY

THIS PAGE IS BLANK